

Mathematical Sciences Technical Report

A Unified Parallel C Implementation for the maximum clique problem

Lisa R. Thimm, Donald Kreher and Phillip Merkey

Michigan Technological University
Mathematical Sciences Technical Report

05-1

February, 2005

MichiganTech

Department of Mathematical Sciences
Houghton, MI 49931-1295
www.math.mtu.edu

Contents

1	Introduction	3
1.1	Cliques	3
1.1.1	Example: Cliques, Maximal and Maximum Cliques	3
1.2	Maximum Clique Applications	4
1.2.1	Example: Relation between a design and a clique	5
1.3	Parallel Computing	6
2	Background	9
2.1	The Maximum Clique Algorithm	9
2.1.1	Example: MaxClique2	10
2.2	UPC	11
2.2.1	Example: UPC barriers	13
2.2.2	Example: UPC Variable Declaration	14
2.3	The State Space Tree and Load Balancing	14
3	The Stealstack	17
3.1	Theoretical Overview	17
3.2	MaxClique	18
3.2.1	Example: Pushing nodes on the stack	22
4	Implementation Details	23
4.1	Data Structures and Variable Declarations	23
4.2	Function Contents	23
4.3	Experimental Results	24
4.4	Conclusions and Future Work	28

Chapter 1

Introduction

In this report we discuss a parallel programming method for solving the maximum clique problem. Beginning with a serial algorithm and program for finding a maximum clique in a graph we developed a parallel algorithm and program to solve the same problem. The serial and parallel algorithms, the parallel program, and the programming language used for the implementation are all discussed. During the implementation process a work distribution problem was encountered, and the solution we used, the Stealstack, is discussed in detail.

1.1 Cliques

Let $G = (V, E)$ be an undirected graph with no repeated edges and no loops. A *clique* in G is a subgraph in which all vertices are pairwise adjacent. Thus a clique is a complete subgraph. A clique is *maximal* if there is no vertex that can extend the clique. In other words, there are no vertices adjacent to every vertex of the clique that are not already in the clique. A clique of the largest possible size in the graph is called *maximum*. A graph may have more than one maximum clique. Let $S \subseteq V$ be a subset of the vertices. If the subgraph of G induced by the vertices in S is a clique, we say that S forms a clique.

1.1.1 Example: Cliques, Maximal and Maximum Cliques

In the graph displayed in Figure 1.1 the vertices in $\{1,2\}$ form a *clique*. Because vertex 3 can extend this clique, $\{1,2\}$ is not maximal. The vertices in $\{1,2,3\}$ form a *maximal* clique. Vertices $\{1,2,4,5\}$ form a *maximum* clique, the only one in this graph. A list of all 22 cliques in this graph follows:

$\{\}$	$\{3\}$	$\{1,2\}$	$\{2,3\}$	$\{1,2,3\}$	$\{2,4,5\}$
$\{0\}$	$\{4\}$	$\{1,3\}$	$\{2,4\}$	$\{1,2,4\}$	$\{1,2,4,5\}$
$\{1\}$	$\{5\}$	$\{1,4\}$	$\{2,5\}$	$\{1,2,5\}$	
$\{2\}$	$\{0,1\}$	$\{1,5\}$	$\{4,5\}$	$\{1,4,5\}$	




Figure 1.1: A sample graph

1.2 Maximum Clique Applications

The problem of finding a maximum clique in a graph is of interest because many combinatorial problems can be translated into the maximum clique problem. As a simple example consider a maximum independent set in a graph G . An *independent set* is a set of vertices such that no two are adjacent. The size of an independent set is the number of vertices in the set, and a *maximum independent set* is one of the largest size in the graph. Finding a maximum clique in a graph G , therefore, is equivalent to finding a maximum independent set in the the complement of G .

As another example consider $t - (v, k, \lambda)$ designs, with $\lambda = 1$. Such a design is a set of *blocks*, each consisting of k elements of a v -element set, V , such that every t -element subset of V appears in exactly λ , or in this case one, of the blocks. The number of blocks in the design, then, is

$$b = \frac{\binom{v}{t}}{\binom{k}{t}} \lambda.$$

Let A be an array with the columns indexed by all t -subsets of V and the rows indexed by all k -subsets of V . An entry in A is 1 if and only if the t -subset is contained in the k -subset, and 0 otherwise. The rows of A corresponding to a design would be disjoint; their dot product would be zero. Let G be a graph with vertices corresponding to the rows of A . Two vertices are adjacent in G if and only if their corresponding rows are disjoint. Finding a clique of size b in the graph is equivalent to finding the corresponding design. If the size of a maximum clique is less than b , no such design exists. See Example 1.2.1.

An n -factor of a graph $G = (V, E)$ is a spanning subgraph of G in which every vertex has degree n . A *graph factorization* is a decomposition of a graph into edge-disjoint spanning subgraphs $H_0 \cup H_1 \cup \dots \cup H_k$, where each H_i is an n -factor of G . The graph in Figure 1.2 represents a 1-factorization of the cube and the graph in Figure 1.3 represents a 2-factorization of K_9 , the complete graph on nine vertices.

Maximum cliques can be used to find k -factorizations of a graph $G = (V, E)$. Begin by finding all k -factors of G . Create a graph G' where each vertex of G' corresponds to a k -factor in G . Two vertices in G' are adjacent if the corresponding k -factors in G are vertex disjoint. A maximum clique in G' of size $2|E|/nk$ corresponds to a k -factorization of the graph G . If the size of a maximum clique in G' is smaller than $2|E|/nk$ no k -factorization of G exists.





Figure 1.2: A 1-factorization of the cube

Figure 1.3: A 2-factorization of K_9

1.2.1 Example: Relation between a design and a clique

Consider the design with parameters $2 - (7, 3, 1)$. Let $V = \{0, 1, 2, 3, 4, 5, 6\}$. Figure 1.4 gives the matrix A with rows indexed by all triples of elements of V and columns indexed by all pairs of elements of V . The first column of A gives a labeling of the rows that corresponds to the vertices of the graph. Two vertices x_i and x_j are adjacent if and only if the dot product of their corresponding rows in A is zero.

The graph produced by this method is regular of degree 22. The size of a maximum clique in the graph is seven. Because the number of blocks needed to form a $2 - (7, 3, 1)$ design is seven, a maximum clique from the graph can be translated into a set of blocks forming the desired design. One maximum clique is formed by the vertices

$$\{x_1, x_8, x_{12}, x_{16}, x_{24}, x_{26}, x_{32}\}.$$

The resulting design consists of the blocks

$$\{0, 1, 3\}, \{0, 2, 6\}, \{0, 4, 5\}, \{1, 2, 4\}, \{1, 5, 6\}, \{2, 3, 5\}, \{3, 4, 6\}.$$

Another maximum clique, formed by vertices

$$\{x_3, x_6, x_{11}, x_{15}, x_{23}, x_{30}, x_{31}\},$$

produces a design with blocks

$$\{0, 1, 5\}, \{0, 2, 4\}, \{0, 3, 6\}, \{1, 2, 3\}, \{1, 4, 6\}, \{2, 5, 6\}, \{3, 4, 5\}.$$

		01	02	03	04	05	06	12	13	14	15	16	23	24	25	26	34	35	36	45	46	56
x_0	012	1	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
x_1	013	1	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0
x_2	014	1	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0
x_3	015	1	0	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0
x_4	016	1	0	0	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0
x_5	023	0	1	1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0
x_6	024	0	1	0	1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0
x_7	025	0	1	0	0	1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
x_8	026	0	1	0	0	0	1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0
x_9	034	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0
x_{10}	035	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0
x_{11}	036	0	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0
x_{12}	045	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0
x_{13}	046	0	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0
x_{14}	056	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1
x_{15}	123	0	0	0	0	0	0	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0
x_{16}	124	0	0	0	0	0	0	1	0	1	0	0	0	1	0	0	0	0	0	0	0	0
x_{17}	125	0	0	0	0	0	0	1	0	0	1	0	0	0	1	0	0	0	0	0	0	0
x_{18}	126	0	0	0	0	0	0	1	0	0	0	1	0	0	0	1	0	0	0	0	0	0
x_{19}	134	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	1	0	0	0	0	0
x_{20}	135	0	0	0	0	0	0	0	1	0	1	0	0	0	0	0	0	1	0	0	0	0
x_{21}	136	0	0	0	0	0	0	0	1	0	0	1	0	0	0	0	0	0	1	0	0	0
x_{22}	145	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	1	0	0
x_{23}	146	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0	1	0
x_{24}	156	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0	1
x_{25}	234	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	1	0	0	0	0	0
x_{26}	235	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	1	0	0	0	0
x_{27}	236	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1	0	0	1	0	0	0
x_{28}	245	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	1	0	0
x_{29}	246	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0	1	0
x_{30}	256	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	1
x_{31}	345	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	1	0	1	0
x_{32}	346	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	1	0
x_{33}	356	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	1
x_{34}	456	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1

Figure 1.4: Matrix A

1.3 Parallel Computing

Traditionally computers have been composed of one central processing unit, CPU, simply called a processor, which can process one instruction at a time. This kind of machine is called a serial machine, and programs written for it are serial programs. Parallel computing, in contrast to serial computing, takes advantage of multiple processes working concurrently to solve the same problem. These processes can run on either a single or multiple processors, which in turn might be part of a single machine, may be located on many individual machines which have been joined into a cluster, or a combination of the two.

Consider 10 integers, $a_0, a_1, b_0, b_1, \dots, e_0, e_1$, where every pair x_0, x_1 needs to be added. A serial machine would require five steps to complete the computation, one for each addition. If the addition were done on a parallel machine with five processes, where each process was assigned a pair of numbers, the computation would only take one step. During that one step each process would add the pair of numbers assigned to it.

In a serial machine the single processor has access to all memory in the machine. Memory in a parallel machine comes in one of two basic forms. In a shared memory machine all processors have access to the




Figure 1.5: Shared and distributed memory machines

same memory. In a distributed memory machine each processor has its own memory. See Figure 1.5 for an illustration of the two memory models, where p denotes a processor and m denotes memory.

One motivation for computing in parallel is the potential speed increase, as seen in the simple addition example above. If one processor can complete a task in five minutes, then perhaps five processors working together can complete the same task in one minute. The speed increase, however, may not be proportional to the number of processors for many reasons. Communication between the processors creates overhead not present in a serial machine. If the problem is not easily broken into equal pieces there may be load balancing issues; some processors may finish early and sit idle while the others continue executing.

Chapter 2

Background

2.1 The Maximum Clique Algorithm

In their book Kreher and Stinson describe a deterministic serial algorithm for finding a maximum clique in a simple graph [1]. Pseudocode for their algorithm, `MAXCLIQUE2(ℓ)`, is given as Algorithm 2.1.1. The parallel algorithm and implementation discussed in this paper are based on Algorithm 2.1.1. The `MAXCLIQUE2(ℓ)` algorithm performs a depth-first search using bounding functions to prune the state space tree and reduce search time.

Let $G = (V, E)$ be a graph on n vertices. Let $V = \{0, 1, \dots, n - 1\}$ and let A_m be the set of vertices in G adjacent to vertex m . `MAXCLIQUE2(ℓ)` begins with a clique, X , of size ℓ with $X = [x_0, x_1, \dots, x_{\ell-1}]$ and computes a set of vertices that can potentially extend X to a clique of size $\ell + 1$. A clique of size k can be written down in $k!$ ways, so a naive algorithm would find the same clique $k!$ times, slowing down runtime. To prevent this an ordering is imposed on the vertices. Different orderings may produce different sets of vertices as the first maximum clique found. We use the natural ordering $0 < 1 < \dots < n - 1$. The `MAXCLIQUE2(ℓ)` algorithm considers as possible vertices only those that are greater than the largest vertex in the current clique X . Proceeding in this manner will produce a clique X such that $x_0 < x_1 < \dots < x_{\ell-1}$. Let B_m be the set of vertices greater than vertex m . The *extension set*, the set of vertices that can potentially extend the clique X of size ℓ to a clique of size $\ell + 1$, includes all vertices adjacent to every vertex of X that are also greater than $x_{\ell-1}$. Call this extension set C_ℓ . Then $C_\ell = A_{x_{\ell-1}} \cap B_{x_{\ell-1}} \cap C_{\ell-1}$ is the set of possibilities for x_ℓ .

Our goal is to find a maximum clique, so there is no need to record every clique found by the algorithm. Instead, only the largest clique currently found is saved. Every clique found is compared to the current largest clique, `OptClique`, and if a larger clique is found it becomes the new `OptClique`. The variable `OptSize` is the size of the current optimal clique.

To reduce the number of nodes in the state space tree a bounding function, `BOUND(X)`, can be included in the algorithm. The simplest bounding function adds the number of vertices in the extension set C_ℓ to the size of the current clique X of size ℓ . If this sum is less than or equal to the size of the current optimal clique, it is impossible to form a larger clique and there is no need to continue processing X . This is called the {size bound}. `MAXCLIQUE2(ℓ)`, complete with bounding function, is given in Algorithm 2.1.1.

Algorithm 2.1.1: MAXCLIQUE2(ℓ)

```

external BOUND()
global  $A_m, B_m, C_m$  ( $m = 0, 1, \dots, n - 1$ )
if  $\ell > OptSize$ 
  then  $\begin{cases} OptSize \leftarrow \ell \\ OptClique \leftarrow [x_0, \dots, x_{\ell-1}] \end{cases}$ 
if  $\ell = 0$ 
  then  $C_\ell \leftarrow V$ 
  else  $C_\ell \leftarrow A_{x_{\ell-1}} \cap B_{x_{\ell-1}} \cap C_{\ell-1}$ 
 $M \leftarrow B([x_0, \dots, x_{\ell-1}])$ 
for each  $x \in C_\ell$ 
  do  $\begin{cases} \text{if } M \leq OptSize \\ \text{then return } () \\ x_\ell \leftarrow x \\ MAXCLIQUE2(\ell + 1) \end{cases}$ 
main
 $OptSize \leftarrow 0$ 
MAXCLIQUE2(0)
output ( $OptClique$ )

```

Other bounding functions produce better results than the simple size bound described above. Kreher and Stinson examined two additional bonding functions, the *greedy bound* and the *sampling bound*. Between these three bounds Kreher and Stinson found the greedy bound to yield the best pruning of the state space tree [1]. At each step the greedy bound greedily colors the vertices of the extension set C_ℓ such that no two adjacent vertices receive the same color. The number of colors needed gives an upper bound on the number of vertices that can be added to the current clique.

The third bound discussed by Kreher and Stinson is the sampling bound. It performs a greedy coloring of all the vertices in the graph before the search begins. At each step the number of colors used on the vertices of C_ℓ is added to the size of the current clique.

2.1.1 Example: MaxClique2

Consider the graph displayed in Figure 2.1. If the sampling bound were being used it would perform an initial greedy coloring of the graph. The vertices $\{0,2,4\}$ would be given color 0, vertices $\{1,3\}$ color 1, vertices $\{5,6\}$ color 2, and vertex 7 would have color 3. At some point $X = [3]$ is the current clique with size $\ell = 1$. Then $A_3 = \{0, 2, 4, 5, 6, 7\}$, $B_3 = \{4, 5, 6, 7\}$, and $C_2 = \{4, 5, 6, 7\}$. We can compare the results of the size, greedy, and sampling bounds. Using the size bound we know that X can be extended into a clique of size at most $\ell + |C_\ell| = 5$. The greedy bound would color the vertices of C_2 as follows: vertex 4 is color 0, vertices 5 and 6 are color 1, and vertex 7 is color 2. The upper bound would therefore be four.

The sampling bound used three colors to color the vertices in C_2 , and would therefore also give a bound of four.

In the next step $X = [3, 4]$ is the current clique with size $\ell = 2$. Then $A_4 = \{3, 5, 6, 7\}$, $B_4 = \{5, 6, 7\}$ and $C_3 = A_4 \cap B_4 \cap C_2 = \{5, 6, 7\}$. The size bound still says that X can be extended into a clique of size at most five. Using the greedy bound, which would color the vertices of C_3 with two colors, we still have an upper limit of four vertices. The sampling bound colored the vertices in C_3 with two colors, giving an upper bound of four.




Figure 2.1: A sample graph

2.2 UPC

Before presenting the parallel algorithm for solving the maximum clique problem we will discuss the programming language used for the implementation. *Unified Parallel C*, or UPC, is an extension of the ANSI C programming language [4]. UPC programs follow the single program, multiple data (SPMD) programming model [3]. The user writes only one piece of code. During execution many copies of this single program are run simultaneously. Each instance of the program is called a thread. There are two constants in UPC that help manage these multiple threads: `THREADS` is the number of threads running and `MYTHREAD` is a unique constant from the set $(0, 1, \dots, \text{THREADS} - 1)$ which provides the thread with an identity.

Unlike the master/slave programming model or the fork/join model of OpenMP, there is no single controlling process that invokes parallel execution within the execution of the program. With UPC all threads run the program from start to finish. Within the execution each instance can use its value of `MYTHREAD` to make decisions within its own execution path. For example, consider the following code.

```
if (MYTHREAD == 0)
    readfile();
```

When each thread reaches this section of code it will consider the `if` statement. Only the thread with `MYTHREAD` equal to zero executes the function `readfile()`. All other threads skip that statement.

Because all threads execute on their own clocks and take different execution paths through the program there is no notion of instruction-by-instruction synchronization. All synchronization must be done explicitly. There are several ways to do this, but the principle way is to use a `UPC_BARRIER()`. Once a thread reaches a

barrier it cannot continue with execution until all other threads have reached the barrier. See Example 2.2.1 for an example of the effects of a `UPC_BARRIER()` in program execution.

UPC is based on a partitioned shared memory model. This is a partition of available memory that aligns with the threads. Partitioned shared memory requires that we declare *shared variables*, variables accessible to all threads, as `shared`. Any variable declared without the keyword `shared` is a *private variable*, and each thread has its own instance of the variable. See Figure 2.2 for a picture of the partitioned memory. The space above a thread represents the memory to which that thread has *affinity*. Each thread can make memory references to its part of memory without interfering with the memory being acted upon by other threads.




Figure 2.2: UPC Memory allocation

Shared variables must have global scope. Consider the following declarations.

```
shared int A;
shared int x[4];
shared int y[20*THREADS];
```

The declaration `shared int A;` means that `A` is an integer variable in the shared space. By convention single variables have affinity to thread zero. The variable `x` is an array that is spread across the shared spaces such that the affinity of `x[i]` is thread T_t where $t \equiv i \pmod{\text{THREADS}}$. Similarly `y` is an array of length `20*THREADS` with affinity assigned as above. Therefore each thread has affinity to 20 elements of array `y`. Example 2.2.2 illustrates shared and private variable declarations.

One way to protect segments of code, and consequently access to shared variables, is to use locks. A *lock* is an object that can be held by only one thread at a time. Before accessing a certain area of shared memory a thread can be required to obtain a lock associated with that area. When the thread is finished it releases the lock. Until the lock is released any other thread attempting to acquire it must wait. Provided that all sections of code that read from or write to a specific area of memory are protected by the same lock, threads are prevented from reading memory that is concurrently being written to, and vice versa.

One of the main benefits of UPC as a parallel language is its ease of use. Because it is based on the C programming language, almost all syntax is familiar to C programmers. MPI, a library that handles explicit message passing, allows processes to send each other copies of the variables they have under their control. UPC, in contrast, achieves communication by references to shared variables, sometimes called implicit message passing. Users can write assignment statements without knowing which thread controls each variable.

2.2.1 Example: UPC barriers

Consider the following UPC code, compiled on four threads, and the resulting output. Without some synchronization the output depends entirely on the speed of each thread.

```
int main()
{
    printf("Thread %d likes bread and butter.\n", MYTHREAD);
    printf("Thread %d likes toast and jam.\n", MYTHREAD);
    return 0;
}
```

```
Thread 0 likes bread and butter.
Thread 2 likes bread and butter.
Thread 1 likes bread and butter.
Thread 0 likes toast and jam.
Thread 2 likes toast and jam.
Thread 3 likes bread and butter.
Thread 3 likes toast and jam.
Thread 1 likes toast and jam.
```

Now consider the modified code below. When a thread reaches a barrier in its execution it cannot continue until all threads have reached the barrier. Although before and after the barrier the order of execution is based on the speed of each thread, no instructions before or after the barrier can be interchanged. In other words, all threads must make their fondness for bread and butter known before any can mention toast and jam.

```
int main()
{
    printf("Thread %d likes bread and butter.\n", MYTHREAD);
    upc_barrier();
    printf("Thread %d likes toast and jam.\n", MYTHREAD);
    return 0;
}
```

```
Thread 0 likes bread and butter.
Thread 2 likes bread and butter.
Thread 1 likes bread and butter.
```

```

Thread 3 likes bread and butter.
Thread 0 likes toast and jam.
Thread 2 likes toast and jam.
Thread 3 likes toast and jam.
Thread 1 likes toast and jam.

```

2.2.2 Example: UPC Variable Declaration

The following variable declarations, compiled to run on four threads, result in the placement shown in Figure 2.2. Notice that a *blocking factor* has been added to the shared declaration of array y . A blocking factor of x means that x elements of the array will be placed in each thread's shared memory, wrapping around if the size of the array is larger than the number of threads. The default blocking factor is 1.

```

shared int AB;
shared int x[5];

shared 2 int y[10];

int a;

```




Figure 2.3: UPC Declarations

2.3 The State Space Tree and Load Balancing

To obtain good performance from parallel computation it only makes sense that all processors, or UPC threads, should be busy during the entire execution of the program. Idle threads contribute little to the speed increases looked for in a parallel implementation. Not all algorithms are capable of being broken into smaller pieces and are therefore not suitable for parallel computation. Other algorithms may be granular enough yet still suffer from load balancing problems for various reasons. To obtain an efficient algorithm the workload between the threads must be balanced.

To better understand the potential for uneven work distribution in a parallel implementation it helps to consider the *state space tree* of the backtrack algorithm. This tree is essentially a picture of the search

conducted by the algorithm, with each vertex or node on the tree representing a clique considered and analyzed by the algorithm. By considering the state space tree of the complete graph on n vertices we can see that any state space tree is at most n nodes deep, and if only vertices of higher order are considered in the extension set, there are at most 2^n nodes in the tree. As an example consider Figure 2.4, the state space tree of the serial algorithm $\text{MAXCLIQUE2}(\ell)$ for the graph in Figure 1.1 that results if no bounding function is used.




Figure 2.4: State space tree for the graph in Figure 1.1

A naive approach to distributing the work would be to assign the threads branches of the tree beginning with the cliques of size one. Consider again the graph in Figure 1.1 and the corresponding state space tree in Figure 2.4. If, for example, there were three threads searching we could assign the branches beginning with cliques [0] and [3] to thread T_0 , branches beginning with cliques [1] and [4] to thread T_1 , and branches with cliques [2] and [5] to thread T_2 . Let each thread have a private set S such that S is the set of cliques of size one assigned to that thread. Each thread would execute the $\text{MAXCLIQUE2}(\ell)$ algorithm just as in the serial version, with a few small changes. Instead of making the first call to $\text{MAXCLIQUE2}(\ell)$ with the current clique X being the empty clique, each thread would make their first call with the current clique X set to one of the cliques of size one in S . Each thread would repeat this for every element of their set S . Pseudocode for this naive parallel approach is given as Algorithm 2.3.1, with the initial distribution of cliques done modulus the number of threads. Any other simple distribution could also be used. Note that only one thread needs to set the value of OptSize to zero. Note also that each thread must acquire a lock before accessing the shared variables OptSize and OptClique .

Processing the graph in Figure 1.1 using Algorithm 2.3.1 on three threads results in thread T_1 having potentially much more work than either of the other two threads. In general the branches toward the left of the state space tree will generate more work than branches toward the right of the tree, because only vertices of higher order are considered. The pruning, or bounding, functions make the unbalance even more extreme. Partial cliques toward the right of the tree are more likely to have smaller extension sets, and those branches will be pruned sooner.

Although it is easy to say that the branches on the left will usually provide more work, it is impossible

to say exactly which branches will produce more work, and how much more. In Figure 2.4, for example, the branch of the state space tree beginning with the clique [0] terminates immediately, while the branch beginning with the clique [1] results in a maximum clique. This means that any initial division of work cannot guarantee that no thread will become idle while others continue processing. What is needed is a way to redistribute work while processing. We chose to redistribute the work using a stealstack.

Algorithm 2.3.1: PARMAXCLIQUE(ℓ)

```

external BOUND()
Shared OptClique, OptSize
global  $A_m, B_m, C_m, \quad (m = 0, 1, \dots, n - 1)$ 
LOCK()
if  $\ell > \textit{OptSize}$ 
  then  $\begin{cases} \textit{OptSize} \leftarrow \ell \\ \textit{OptClique} \leftarrow [x_0, \dots, x_{\ell-1}] \end{cases}$ 
UNLOCK()
if  $\ell = 0$ 
  then  $C_\ell \leftarrow V$ 
  else  $C_\ell \leftarrow A_{x_{\ell-1}} \cap B_{x_{\ell-1}} \cap C_{\ell-1}$ 
 $M \leftarrow \text{BOUND}([x_0, \dots, x_{\ell-1}])$ 
for each  $x \in C_\ell$ 
  do  $\begin{cases} \text{if } M \leq \textit{OptSize} \\ \quad \text{then return } () \\ x_\ell \leftarrow x \\ \text{PARMAXCLIQUE}(\ell + 1) \end{cases}$ 
main
if MYTHREAD == 0
  then OptSize  $\leftarrow 0$ 
for  $s \leftarrow 0$  to  $n - 1$ 
  do if  $i/\text{THREADS} == \text{MYTHREAD}$ 
  then  $S \leftarrow S \cup s$ 
for each  $s \in S$ 
  do  $\begin{cases} x_0 \leftarrow s \\ C_0 \leftarrow A_s \\ \text{PARMAXCLIQUE}(1) \end{cases}$ 
output (OptClique)

```

Chapter 3

The Stealstack

3.1 Theoretical Overview

The *stealstack* was first implemented by Prins et al. as a means of dynamically balancing the work load of an unbalanced search tree [2]. Each thread has affinity to a stack and an associated lock. The stacks and their locks are declared in shared memory, allowing all threads access to each stack. The nodes on a stack are data structures that correspond to nodes in the state space tree. Each node contains enough information to start a search as a root of a subtree of the complete state space tree. These nodes are divided into two categories. The nodes on the top of the stack are called the *local nodes*, and those on the bottom are the *shared nodes*. The local and shared portions of a stack should not be confused with private and shared memory in UPC. The local nodes are for processing only by the thread with affinity to the specific stack. Shared nodes are free to be stolen by any idle thread. The lock associated with a stack controls access to the shared nodes of the stack. See Figure 3.1.

When a thread anticipates that it will have too much work, a decision that is discussed later, it pushes jobs onto the top of the local portion of its stack. A thread can also release nodes from the local to shared portion of its stack, freeing them to be stolen. When the thread finishes processing its current work, it attempts to pop a job off the top of its local stack. If the local stack is empty the thread will attempt to move nodes from the shared portion of its stack to the (empty) local portion. The thread first acquires its stack lock, the lock associated with its particular stack. Once the lock has been acquired the thread checks to see that there are nodes in the shared portion, and if so moves some of them to its local portion and releases the lock. If the thread acquires the lock and discovers that its shared portion is also empty, it releases the lock and attempts to steal nodes from another thread.

The idle thread begins the stealing process by examining the other threads' stacks, searching for one with nodes in the shared portion of its stack. When the idle thread finds such a stack, it attempts to acquire the lock associated with it. Once the lock is acquired the idle thread first makes sure that there are still nodes available, and if so moves them to the local portion of its own stack, then releases the lock. Any number of polling methods can be used. The simplest way for thread T_i to find a thread with work to steal is to begin with thread $T_{i+1 \pmod{\text{THREADS}}}$ and loop through all threads sequentially. The polling need not be sequential, however. It could follow a random pattern, a predetermined permutation of the threads, or

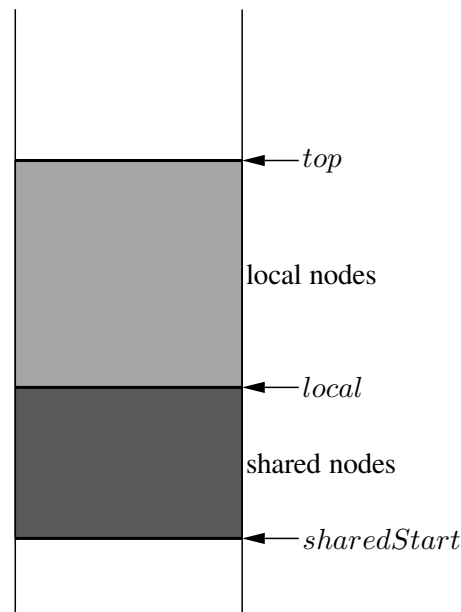


Figure 3.1: A stack

any cycle the user determines will produced better results. The number of times a thread polls is also open to optimization by the user.

Having the idle threads look for work is an advantageous feature of the stealstack, and why it was chosen to solve this problem. Some load-balancing models have busy threads pause their computation to find idle threads to give their work to.

Figure 3.2 summarizes the functions associated with the stealstack.

3.2 MaxClique

The stealstack now needs to be incorporated into the $\text{MAXCLIQUE2}(\ell)$ algorithm. Each node, or job, on the stack needs to contain three elements: a clique, its size, and its extension set. Let $temp$ be a node. We will denote node $temp$'s clique by $temp.X$, the size of the clique by $temp.L$, and the extension set by $temp.C$. The node may also be denoted by $temp_{X,C,L}$.

At some point the thread needs to decide if it has too much work, and if so, what work it will put on its stack. Let $\text{TMW}()$ be a function that returns **true** if the thread has too much work, and **false** otherwise. A particular implementation of $\text{TMW}()$ is discussed Section 4.2. In $\text{MAXCLIQUE2}(\ell)$ each element of the extension set C_ℓ is included, in turn, in the current clique, and the search continues from there. In the parallel version, Algorithm 3.2.1, if the thread determines that it has too much work it includes the first element of C_ℓ in the current clique. Before processing this new clique the algorithm includes each additional element of C_ℓ in the previous clique and pushes the job, the new clique and its size and extension set, onto its stack. At this point the thread also checks the number of jobs in the local portion of its stack. The function $\text{FULL}()$ returns **true** if there are more than enough jobs in the local portion, and **false** otherwise. An implementation of

Function Name	Action Taken	Return Value
PUSH(&someNode)	the work in someNode is pushed on top of the stack	none
POP(&someNode)	the work from the node on the top of the stack is put into someNode	none
LOCALDEPTH()		number of nodes in local part of stack
RELEASE(k)	k nodes are moved from local to shared part of stack	none
ACQUIRE(k)	k nodes are moved from shared to local part of stack	TRUE if k nodes are moved from shared to local part of stack, else FALSE
STEAL(t, k)	k nodes moved from bottom of thread t 's shared portion to top of local portion	TRUE if k nodes are taken from thread t 's stack, else FALSE
FINDWORK(k)		number of a thread with at least k nodes in shared part, -1 if no such thread

Figure 3.2: Stealstack Functions

FULL() is discussed in Section 4.2. If FULL()= **true** then the thread releases some of the nodes from the local to the shared portion of its stack, freeing them to be stolen. It then proceeds with the new current clique. See Example 3.2.1.

Another property of the algorithm to note is the initial distribution of the work. In the serial version the algorithm essentially starts with nothing. It has no current clique, so its first action is to create C_0 as the set of all vertices. Each of these will at some point become the first element of a clique. In the parallel version it makes sense to begin by distributing these cliques of size one to all threads. Giving them all work to start on immediately saves time that would be lost if only one thread began working and all others had to wait for that thread to generate enough work for them to steal.

Algorithm 3.2.1: MAXCLIQUE(ℓ)

```

external { BOUND(), TMW()
           PUSH( $c$ ), FULL()
           RELEASE( $k$ )
global  $A_\ell, B_\ell, C_\ell$ 
LOCK()
if  $\ell > OptSize$ 
  then {  $OptSize \leftarrow \ell$ 
         $OptClique \leftarrow [x_0, \dots, x_{\ell-1}]$ 
UNLOCK()
if  $\ell = 0$ 
  then  $C_\ell \leftarrow V$ 
  else  $C_\ell \leftarrow A_{x_{\ell-1}} \cap B_{x_{\ell-1}} \cap C_{\ell-1}$ 
 $M \leftarrow \text{BOUND}([x_0, \dots, x_{\ell-1}])$ 
if TMW()
  then { if  $M \leq OptSize$ 
        then return ()
         $x_\ell \leftarrow$  the first  $x \in C_\ell$ 
        for each additional  $x \in C_\ell$ 
          do  $push(x, C_\ell)$ 
          if FULL()
            then RELEASE( $k$ )
          MAXCLIQUE( $\ell + 1$ )
        else for each  $x \in C_\ell$ 
          do { if  $M \leq OptSize$ 
              then return ()
               $x_l \leftarrow x$ 
              MAXCLIQUE( $\ell + 1$ )

```

Algorithm 3.2.1 is a modified version of Algorithm 2.1.1. It differs in that the conditional TMW(), the subsequent addition of nodes to the stack, and the conditional FULL() and following RELEASE(k) have been added. Although this now allows threads to add work to their stack, and even release work for stealing, it does not allow a thread to remove work from its own stack, or steal from other threads. An algorithm for enabling a thread to empty its stack will be addressed first. Consider Algorithm 3.2.2. The initial jobs, all cliques of length one, are distributed among the threads and placed on their local stacks. In the function DFS() a node is removed from the local stack and processed completely. During the processing other nodes may or may not be added to the local stack. When the thread finishes with that portion of the graph it begins on the next node in its local stack. When the local portion is empty the thread attempts to move nodes from

its shared to local portion. The function ACQUIRE(a) returns **true** if a nodes were moved from the shared to local portion, and **false** if there were not a nodes to move. When DFS() completes the thread has no nodes in either its local or shared stack, and is out of work.

Algorithm 3.2.2: DFS()

```

external { MAXCLIQUE()
            { POP( $c$ )
repeat
  while LOCALDEPTH() > 0
    do {  $temp \leftarrow POP()$ 
        {  $X \leftarrow temp.X$ 
          {  $C_{\ell+1} \leftarrow temp.C$ 
            { MAXCLIQUE( $temp.L$ )
           $moreWork \leftarrow ACQUIRE(a)$ 
        until  $moreWork = FALSE$ 
main
  for each  $x < n, x \equiv MYTHREAD(\text{MOD } n)$ 
    do PUSH( $c_{x,1}, B_x$ )
   $OptSize \leftarrow 0$ 
  BARRIER()
  DFS()
output ( $OptClique$ )

```

Now what is needed to balance the load is an algorithm that incorporates stealing. Consider Algorithm 3.2.3. The function DFS() empties both the local and shared portion of a thread's stack. Once that happens the thread needs to find work to steal, and begins searching. The function FINDWORK() polls the other threads using whatever method is decided upon, and returns the number of a thread with enough work to steal. If no other thread has enough nodes in its shared portion FINDWORK() returns -1 and the thread exits. Once a thread with potential work is found the nodes must be moved from its shared stack to the local stack of the searching thread. The function STEAL(victimId, k) returns **true** if k nodes are successfully stolen from thread $T_{victimId}$ and **false** otherwise. Because another thread may have moved the nodes to its own stack between the time the searching thread locates the busy thread and attempts to move the nodes, merely finding a thread with enough work is not a guarantee that a thread will be able to steal the work.

Algorithm 3.2.3: EXPLOREFOREST()

```

external { DFS(), FINDWORK()
            STEAL( $t, k$ )
          }
DFS()
 $victimId \leftarrow$  FINDWORK()
while  $victimId \neq -1$ 
  do { if STEAL( $victimId, k$ )
        then DFS()
         $victimId \leftarrow$  FINDWORK()
      }
main
for each  $x < n, x \equiv$  MYTHREAD( MOD  $n$ )
  do PUSH( $c_{x,\ell}, B_x$ )
   $OptSize \leftarrow 0$ 
  BARRIER()
  EXPLOREFOREST()
output ( $OptClique$ )

```


3.2.1 Example: Pushing nodes on the stack

Figure 3.3: A sample graph

Consider the graph displayed in Figure 3.3 Let $TMW()$ return **true** if and only if $|C_\ell| > 3$ and let $X = [2, 3]$ be the current clique. Then $\ell = 2$ and $C_3 = \{4, 5, 6, 7\}$. The new clique is then $X = [2, 3, 4]$ and the nodes

$$\ell = 3 \quad X = [2, 3, 5] \quad C = \{4, 5, 6, 7\}$$

$$\ell = 3 \quad X = [2, 3, 6] \quad C = \{4, 5, 6, 7\}$$

$$\ell = 3 \quad X = [2, 3, 7] \quad C = \{4, 5, 6, 7\}$$

are pushed on top of the stack.

Chapter 4

Implementation Details

4.1 Data Structures and Variable Declarations

The main data structure used in the program is a set, represented by a bit vector. Element i is in the set if and only if the i th bit is 1. The graph, then, is stored as an array of sets, $[A_j : j = 0, 1, \dots, n - 1]$. The j th set in the array represents the vertices in G adjacent to vertex j . Current cliques and the current optimal clique are also stored as sets. Using bit vectors to represent the sets results in a smaller amount of memory needed to store the data than if data structures like linked lists or arrays of integers had been used. It also facilitates basic set operations; intersections and unions of sets are quick and simple to compute. Because UPC sometimes has difficulty dealing with dynamically allocated variables everything is declared statically before compilation. This means an upper bound on the number of vertices in the graph needs to be known beforehand.

Because the data structure used to store the graph information is relatively small, each thread makes a copy of the graph in its own private memory. The program references the graph frequently. Thus, placing a copy in memory that a thread has affinity to will speed up program execution.

Other than the stealstacks the only data placed in shared memory are the current optimal clique, `OptClique`, its size, `OptSize`, and the lock that regulates their access. Because all threads access these variables there is no benefit to giving them affinity to any particular thread, and they are declared simply as shared.

Distributing the initial cliques of length one, represented by the children of the root, modulus the thread number, divides more evenly the heavier branches among the threads than simply partitioning the branches. This way each of the threads will hopefully begin with one of the heavier branches from the left side of the tree.

4.2 Function Contents

The function `TMW()` determines whether or not the thread will push some of its jobs onto its stealstack. Because `TMW()` is called with each recursive call of `MAXCLIQUE()` it must run quickly. We chose to use the size of the extension set as the criteria for determining the amount of work the current clique will produce. A large extension set has the potential for creating much more work than a small set. Because the

sets are all stored as bit vectors a simple mask and lookup algorithm for finding the size of the sets is not overly time-consuming.

The function `FULL()` determines if nodes will be released from the local to the shared portion of a thread's stealstack. As with the function `TMW()` the function `FULL()` is called with every recursive call of `MAXCLIQUE()` and must therefore run quickly. We chose to use the number of nodes in the local portion to determine if nodes should be released. If `LOCALDEPTH()`, a stealstack function, returns a value above a set limit, the thread will release the requisite number of nodes to its shared portion.

The number of nodes to release to the shared stack, acquire from the shared stack, and steal from another thread are all open to optimization. Stealing several nodes at a time cuts down on overhead. Without care, however, a problem may arise. The number of nodes released at a time, acquired at a time, and stolen at a time must be coordinated. If they are not it may happen that a thread has a number of nodes in the shared portion of its stealstack, but there are too few to steal and too few to acquire. In that case the nodes would go unprocessed and the program might return a clique that is not maximum. One way to ensure that no nodes are left in the shared stack is to make the number of nodes stolen at a time equal to the number of nodes released and the number acquired, say k . Then the number of nodes in the shared stack at any time would be a multiple of k and there would never be too few to be stolen or acquired. Another option is to acquire only one node at a time. With this option the number released and the number stolen at a time do not need to be equal. This allows the thread to release a number of nodes proportional to the number in its local stack.

4.3 Experimental Results

The motivation for a parallel implementation of the maximum clique finding algorithm was an increase in the speed of execution. To determine the validity of this presumption we compared the time of execution of the serial program to that of the parallel program. We first determined the correctness of the parallel program by running multiple trials on graphs up to 200 vertices and comparing the results to the serial program. We then generated random graphs on 250, 500, 1000, and 2000 vertices with densities of .20, .35, and .50. A density of k means that each possible edge is in the graph with probability k . This resulted in 12 random graphs. We first ran the serial program on each of these graphs. To obtain comparable results the serial program, originally a C program, was compiled as a UPC program, then run on one thread. Had the serial program been compiled as a C program, with a C compiler, the execution times may have been slightly faster. We then ran the parallel program, using the greedy bound, on the graphs, using varying numbers of threads: 2, 10, 25, and 40 threads.

All trials were run on a Cray T3E using a non-conforming compiler. This was done for several reasons. First, the T3E is a fast platform, and allowed larger graphs with higher densities to be searched than would otherwise have been possible. The version of UPC running on the T3E, though not compliant, is stable. The number and types of UPC functions used in the program are small and simple enough that using the non-compliant compiler was not an issue. The program can easily be brought up to compliance and run on other machines.

The results are given in Tables 4.1, 4.2, and 4.3. The first number in each entry is the time in seconds

Table 4.1: Density .20: Number of vertices vs number of threads: Time in seconds and total number of nodes processed

n	Number of threads				
	1	2	10	25	40
250	1.392	0.840	0.271	0.272	0.308
	4,306	4,320	4,292	4,138	4,144
500	15.851	5.006	1.291	0.824	0.724
	21,046	21,301	21,808	22,996	22,225
1000	368.160	76.359	16.279	7.259	5.082
	302,596	302,726	303,252	302,810	302,956
2000		1654.030	337.225	136.812	88.303
		5,066,860	5,084,782	5,062,776	5,062,985

Table 4.2: Density .35: Number of vertices vs number of threads: Time in seconds and total number of nodes processed

n	Number of threads				
	1	2	10	25	40
250	7.593	5.015	1.184	0.650	0.590
	25,670	25,939	27,309	27,819	29,329
500	209.802	83.019	17.397	7.754	5.419
	345,207	347,616	356,780	383,842	410,185
1000	13,368.565	3,505.992	699.551	277.631	178.432
	12,449,987	12,559,193	12,520,920	12,309,827	12,627,817
2000				13,316.92	8,488.096
				402823073	411,591,405

for execution of the program. The second number is the total number of nodes processed during execution. The columns are labeled with the number of threads, with one thread representing the serial program. The rows are labeled with the number of vertices, n , in the graph.

The table in Figure 4.1, for example, gives the results for all trial graphs of density .20. Looking at any of the rows it is easy to see a dramatic increase in speed of execution as the number of threads increases. One might expect a speed increase directly proportional to the number of threads used. However, the tables show this to be false. Increasing the number of threads from one to two does not cut the time in half, nor does increasing from one to ten cut the time by a factor of 10. Managing the stealstacks, waiting at barriers and for locks, for example, all produce overhead not present in the serial version. This overhead is one of the factors that prevent the expected increase in speed.

To better see the speed increase brought about by a greater number of threads, consider the graphs in Figures 4.1, 4.2, and 4.3. Each graph shows the completion time, in seconds, for one of the randomly generated graphs, depending on the number of threads utilized. Note that although all three graphs have the

Table 4.3: Density .50: Number of vertices vs number of threads: Time in seconds and total number of nodes processed

n	Number of threads				
	1	2	10	25	40
250	87.435 292,600	61.691 288,577	12.654 290,839	5.435 291,177	3.702 293,026
500		3,488.164 12,924,435	688.334 12,731,349	276.494 12,751,478	173.756 12,771,695
1000					18,313.697 1,008,807,821
2000					




Figure 4.1: Density .20, 1000 vertices: Time in seconds vs Number of threads

same shape, the scales are quite different.

To illustrate that sharing was having a significant effect on the load balancing and efficiency of the program several of the random graphs were processed again, without allowing any stealing. The overhead of the stealstacks was still present, but each thread worked only on the nodes in its own stack. The results of two of these trial runs, both on 500 vertices with density .35, are given in Tables 4.4 and 4.5. Each thread reported the number of nodes it processed, and from that the standard deviation was calculated. Looking at either table it is obvious that allowing threads to steal work from each other dramatically effects the load balance.

Also notable in Tables 4.1, 4.2, and 4.3 is the effect of the number of threads on the total number of nodes searched. It might be expected that the number of nodes searched would decrease, when it actually either remains about the same or even increases.




Figure 4.2: Density .35, 1000 vertices: Time in seconds vs Number of threads




Figure 4.3: Density .50, 250 vertices: Time in seconds vs Number of threads

Table 4.4: 500 vertices, density .35, 10 threads

	total num nodes	std dev	time
sharing	356780	247	17 sec
no sharing	335061	1,843	78 sec

Table 4.5: 500 vertices, density .35, 25 threads

	total num nodes	std dev	time
sharing	383842	188	8 sec
no sharing	327145	1,604	37 sec

Several entries in Tables 4.1, 4.2, and 4.3, especially in Table 4.3 are empty. The time to completion for those graphs with the particular number of threads was too long.

4.4 Conclusions and Future Work

The ideas used to parallelize the maximum clique algorithm can be applied to other algorithms. Any backtracking algorithm can use stealstacks to balance the workload between the threads. A parallel version of Stinson and Kreher's Exact Cover algorithm would be useful [1].

The implementation discussed in this paper utilizes the parallelism for moving through the state space tree. Any work done at a node on the tree is done by a single processor, and is therefore serial. One of the interesting side effects of this method is that, when compared to the serial version, the parallel program visits more nodes in the state space tree. There is a simple reason for what may at first seem counterintuitive. As the serial program traverses the state space tree it increases the size of the optimal clique. As it moves from left to right in the tree it prunes more branches because it has a larger optimal clique to prune against. In the parallel version several processors begin moving down the tree at the same time, and so for a while very little pruning happens because all processors are examining cliques of approximately the same size. The branches on the right side of the state space tree are not being compared to the ultimate largest clique on the left side, like with the serial program, but rather with a clique at about the same level. With less pruning happening more nodes are processed, meaning that the parallel version is actually doing more work. It is only because the work is being divided between several processors that the parallel program is able to complete sooner than the serial.

Moving the parallelism from the movement through the state space tree to the nodes of the tree may result in less work being done. Several of the threads could be designated as *overthreads*, each having a number of threads assigned to it. The overthreads would traverse the tree just like all threads do in the current parallel implementation, but at each node the *underthreads* would perform the bounding work in parallel. If the greedy bound was being used, for example, then the underthreads would perform a greedy coloring algorithm in parallel. Movement through the tree would be slower, but because better pruning is being done fewer nodes may be processed and there may be an overall reduction in work. The problem with this implementation is the potential for the underthreads to sit idle while the overthread decided which node to process next. The time lost by idle threads may or may not make up for the smaller amount of work resulting from improved pruning. A comparative analysis of the two methods would be interesting.

Bibliography

- [1] D.L. Kreher and D.R. Stinson, *Combinatorial algorithms: generation, enumeration and search*, CRC Press, Boca Raton, 1990.
- [2] J.Prins, Jun Huan, B. Bugh, Chau-Wen Tseng and P. Sadayappan, *UPC Implementation of an Unbalanced Tree Search Benchmark*, <http://www.cs.unc.edu/~huan/papers/03-034.pdf>, Univ. of North Carolina at Chapel Hill and Univ. of Maryland at College Park and Ohio State University, 2003.
- [3] M. Flynn, Some Computer Organizations and Their Effectiveness, ITC, **C-21** (1972) 94.
- [4] S. Chauvin, P. Saha, F. Cantonnet, S. Annareddy and T. El-Ghazawi, *UPC Manual*, George Washington University, <http://www2.gwu.edu/~upc/documentation.html>, 2003.